On the Vulnerability of Digital Fingerprinting Systems to Finite Alphabet Collusion Attacks

Jalal Etesami, *Student Member, IEEE,* and Negar Kiyavash *Senior Member, IEEE*

arXiv:1604.08071v1 [cs.CR] 27 Apr 2016

***Abstract*—This paper proposes a novel, non-linear collusion attack on digital fingerprinting systems. The attack is proposed for fingerprinting systems with finite alphabet but can be extended to continuous alphabet. We analyze the error probability of the attack for some classes of proposed random and deterministic schemes and obtain a bound on the number of colluders necessary to correctly estimate the host signal. That is, it requires fewer number of colluders to defeat the fingerprinting scheme. Our simulation results show that our attack is more powerful in practice than predicted by the theoretical bound.**

I. Introduction

DIGITAL fingerprinting schemes have been invented for traitor tracing as a means of copyright protection. To this end, each user is provided with his own individually modified copy of the host signal (i.e., the content that is copyright protected). The modification entails embedding a unique signature, also known as a fingerprint, in the host signal. The fingerprint can later be retrieved by a detector examining an illegal copy of the content in order to implicate the users who took part in the forgery.

A myriad of fingerprint designs and detection procedures have been proposed in the literature. Examples include orthogonal fingerprints [1], regular simplex fingerprints [5], random Gaussian fingerprints [3], [7], and fingerprints with equiangular tight frames (ETF) [10]. The main consideration in design of fingerprints is their rate, i.e., number of copies that could be distributed and their robustness against attacks, in particular collusion attacks.

In a collusion attack, a group of users combine their copies to create a forgery in order to defeat the detection process. For instance, one such type of collusion attack is linear averaging, in which the colluders estimate the host signal as the linear average of their marked copies. This attack is studied in numerous works [2], [4], [6].

A novel collusion attack is introduced in [9] that proves detrimental against both random and deterministic fingerprints as long as they are from a finite alphabet. The entries of random fingerprints are chosen i.i.d. from a finite set according to a distribution while deterministic fingerprints are designed non-randomly often according to some geometric or algebraic property. The analysis in [9] shows that if the size of collusion is of order $O(\log N)$ where N is the length of a real-valued host signal, the attackers can accurately estimate the host signal with high probability. In particular, this attack improves over the previous best-known attack against ETF fingerprints [10], [15], i.e., the linear averaging attack. It was shown that finite alphabet ETF random or deterministic fingerprints can withstand linear averaging attack as long as number of colluders is bounded by $O(N/\log^a M)$ and $O(\sqrt{N})$ respectively, where M denotes the total number of fingerprinted copies [10].

In this work, we analyze the error probability of the aforementioned collusion attack against uniformly symmetric random fingerprints and obtain a bound on the number of colluders necessary to succeed with high probability. Furthermore, we study the performance of the attack against two more probabilistic fingerprints: optimal probabilistic fingerprint (OPF) codes, aka Tardos codes, and column-wise random fingerprints (Tardos-like codes). The OPF codes are introduced in [17] that are ϵ-secure against K pirates and have length $O(K^2 \log 1/\epsilon)$. Our analysis shows that our collusion attack outperforms the linear averaging [12], minority voting [12], and the majority voting attacks [14], the latter two known to be effective collusion attacks against the OPF codes.

We further generalize the application of this collusion attack to continuous alphabet codes, specifically Gaussian fingerprints by quantizing their support set. Our analysis shows a good performance of the generalized attack compared with the uniform linear averaging attack, proved to be optimal among the class of order-statistic collusion attacks, when the colluders are subject to a mean squared distortion constraint [6].

The rest of paper is organized as follows: Section II describes the mathematical model of the fingerprinting problem. In Section III, we present the proposed collusion attack. Section IV, V, and VI examine the attack's effectiveness on both deterministic and random fingerprints and propose lower bounds on the number of colluders that can break the fingerprinting system with high probability. In Section VIII, we use simulations to demonstrate the effectiveness of our attack on two types of fingerprint codes. We present our conclusions in Section IX.

Throughout this paper, we use bold upper-case letters to denote matrices, boldface for vectors, lower-case for scalar values, and calligraphic fonts for sets.

II. Problem Setup

In this section we describe our mathematical setup for fingerprint embedding.

Consider a host signal $\mathbf{s} \in \mathbb{R}^N$ and $M(> N)$ marked copies of the host signal $\mathbf{s}$ distributed among M users. In particular,

J. Etesami is with the Department of Industrial and Enterprising Systems Engineering and the Coordinated Science Laboratory, University of Illinois at Champaign-Urbana, Urbana, Illinois 6180, USA. (e-mail: etesami2@illinois.edu).

N. Kiyavash is with the Department of Industrial and Enterprising Systems Engineering and the Coordinated Science Laboratory, University of Illinois at Champaign-Urbana, Urbana, Illinois 6180, USA. (e-mail: kiyavash@illinois.edu).

the mth user receives the copy,

$$\mathbf{q}_m := \mathbf{f}_m + \mathbf{s}\ , \tag{1}$$

where $\mathbf{f}_m \in \mathbb{R}^N$ is the mth fingerprint. We assume that all fingerprints have equal energy, i.e.,

$$||\mathbf{f}_m||_2^2 = 1\ , m = 1, ..., M, \tag{2}$$

in the case of deterministic construction, and

$$\mathbb{E}[||\mathbf{f}_m||^2] = 1\ , m = 1, ..., M, \tag{3}$$

in the case of random construction.

Let $\mathbf{F}$ denote the $N \times M$ matrix whose columns are the fingerprints $\{\mathbf{f}_m\}_{m=1}^M$, and let $\{\mathbf{e}_m\}_{m=1}^M$ denote the standard basis in $\mathbb{R}^M$, i.e., for any $m \in \{1, ..., M\}$, $e_m(i) = 0$ when $i \neq m$ and 1 otherwise. Then (1) may be written as

$$\mathbf{q}_m = \mathbf{F}\mathbf{e}_m + \mathbf{s}\ . \tag{4}$$

The goal of the fingerprinting process is to prevent users from illegally sharing their copies. Therefore, the fingerprint detector must be able to detect the users or some subset of the users who took part in the forgery.

III. Attack Strategy

In this section we introduce a new type of collusion attack against fingerprints from a finite alphabet.

Consider a fingerprinting matrix $\mathbf{F} = [f_{i,j}]_{N\times M}$ where the entries are chosen deterministically or at random from a finite set $\Xi = \{\xi_1, ..., \xi_l\} \subset \mathbb{R}$, such that $\xi_1 < \xi_2 < ... < \xi_l$. Given that the attackers know $\Xi = \{\xi_1, ..., \xi_l\}$, they can construct the following set:

$$\mathcal{U} = \{u \in \mathbb{R} :\ \exists\ \textit{unique}\ (i, j)\ s.t.\ u = \xi_i - \xi_j\}. \tag{5}$$

Each element of set $\mathcal{U}$ is a real number that corresponds to the difference of a unique ordered pair in $\Xi \times \Xi$. Note that $|\mathcal{U}| \geq 2$ because $\pm(\xi_l - \xi_1)$ are always in $\mathcal{U}$.

Suppose K users collude with the goal of learning at least one of their fingerprints and subsequently obtaining the host signal. Without loss of generality denote their copies by $\{\mathbf{q}_1, ..., \mathbf{q}_K\}$. The attackers form the $N \times K$ matrix $\mathbf{Q} = [\mathbf{q}_1 ... \mathbf{q}_K]$. Furthermore, they choose one of the copies, e.g., $\mathbf{q}_1$ (it will become clear that this choice is immaterial) and compute the $N \times K$ matrix $\mathbf{A}$ whose ith column is given by $\mathbf{q}_1 - \mathbf{q}_i$ or equivalently $\mathbf{f}_1 - \mathbf{f}_i$. Denote the ith row of $\mathbf{A}$ by $\mathbf{a}_i$. Whenever the attackers encounter an element of $\mathbf{A}$, $a_{i,j}$ that belongs to $\mathcal{U}$, they can uniquely determine $(\xi, \xi') \in \Xi \times \Xi$ s.t. $a_{i,j} = \xi - \xi'$. To illustrate how this helps the attackers learn one of the fingerprints, we provide the following example:

Example 1: Assume 4 colluders took part in the forgery. Furthermore, assume $\Xi = \{-1, 0, 1\}$ and let the ith row of $\mathbf{Q}$ be $[s_i, s_i + 1, s_i + 1, s_i - 1]$. Hence, $\mathbf{a}_i = [0, -1, -1, 1]$ and $\mathcal{U} = \{-2, 2\}$. Although no element of $\mathbf{a}_i$ is in $\mathcal{U}$, by computing pairwise differences of elements of $\mathbf{a}_i$, the colluders can learn $f_{i,1}$ as follows: $a_{i,2} - a_{i,4} = (f_{i,1} - f_{i,2}) - (f_{i,1} - f_{i,4}) = f_{i,4} - f_{i,2} = -2$. Therefore, $f_{i,4} = -1, f_{i,2} = 1$ and subsequently from $f_{i,1} - f_{i,4} = 1$, it follows that $f_{i,1} = 0$.

As illustrated by Example 1, looking at entries of $\mathbf{A}$ alone was not enough to learn the fingerprints. Therefore, for each row i of matrix $\mathbf{A}$ the attackers compute the set:

$$\mathcal{A}_i := \{\Delta a_i : \Delta a_i = a_{i,j} - a_{i,k} :\ k, j = 1, ..., K\}. \tag{6}$$

If $\Delta a_i \in \mathcal{U}$, then attackers can learn $f_{i,1}$. This is because $a_{i,j} = f_{i,1} - f_{i,j}$ and $a_{i,k} = f_{i,1} - f_{i,k}$. Thus $\Delta a_i = f_{i,k} - f_{i,j}$ and following the definition of set $\mathcal{U}$, $(f_{i,k}, f_{i,j})$ is uniquely determined. Subsequently $f_{i,1}$ can be computed as

$$f_{i,1} = a_{i,j} + f_{i,j}, \tag{7}$$

Once $f_{i,1}$ is known, all others $f_{i,j}$, can be determined from (7). If no elements of $\mathcal{A}_i$: $\Delta a_i \notin \mathcal{U}$, the attackers estimate $f_{i,1}$ by the first component of $\mathbf{b}_i^* = (b_{i,1}^*, ..., b_{i,K}^*)$, where

$$\mathbf{b}_i^* = \arg \max_{\mathbf{b} \in \mathcal{B}(\mathbf{a}_i)} p(\mathbf{b}), \tag{8}$$

where $p(\mathbf{b})$ denotes the probability of vector $\mathbf{b}$ and $\mathcal{B}(\mathbf{a}_i)$ is given by

$$\mathcal{B}(\mathbf{a}_i) := \Big\{\mathbf{b} = (b_1, ..., b_K) \in \Xi^K :\ b_1 - b_j = a_{i,j},\ j = 1, ..., K\Big\}. \tag{9}$$

That is, the attackers choose the most likely fingerprints that are consistent with the set of observations $\mathcal{A}_i$. If there are several solutions to (8), the attackers choose one of them arbitrarily. The attackers compute $p(\mathbf{b})$ using their side information about the structure of fingerprints. For instance, if they know that the entries of $\mathbf{F}$ were chosen i.i.d. at random according to $\mathbf{p}(\Xi) = (p_{\xi_1}, ..., p_{\xi_l})$, then

$$p(\mathbf{b}) = \prod_{i=1}^{K} p_{b_i}. \tag{10}$$

Recall that attackers could learn $f_{i,1}$ if they could use the Equation (7). Otherwise, their failure probability depends on how $\mathbf{F}$ was generated. Note that by minimizing $|\mathcal{U}|$, the designer increases the likelihood of attackers' failure regardless of the structure of $\mathbf{F}$.

IV. Attack on Equiangular Tight Frame

In this section we study the performance of our attack on a deterministic fingerprint known as equiangular tight frame (ETF) fingerprints from a finite alphabet, proposed in [15]. Although it was shown in [10] that deterministic ETF fingerprints with finite alphabet are robust against the linear averaging attack, namely, if the size of collusion is $O(\sqrt{N})$, then with high probability the focused detector, which performs a binary hypothesis test for each user to decide whether that particular user is guilty [10], is able to detect at least one of the attackers, we will show that these type of fingerprints can not withstand the proposed attack.

A *frame* is a collection of vectors $\{\mathbf{f}_m\}_1^M \in \mathbb{R}^N$ with frame bounds $0 < A \leq B < \infty$ such that

$$A||\mathbf{x}||^2 \leq \sum_{m=1}^{M} |<\mathbf{x}, \mathbf{f}_m>|^2 \leq B||\mathbf{x}||^2,$$

for every $\mathbf{x} \in \mathbb{R}^N$. The frame is tight if $A = B$. An equiangular tight frame (ETF) is a unit norm tight frame with the additional property that there exists a constant c such that

$$| < \mathbf{f}_{m'}, \mathbf{f}_m > | = c, \ \forall \ m \neq m'.$$

The construction in [15] uses a tensor-like combination of a Steiner system's adjacency matrix and a regular simplex.

A Steiner system [8] with parameters r, h, n, written as $S(r, h, n)$, is an n-element set $\mathcal{S}$ together with a set of h-element subsets of $\mathcal{S}$ (called blocks) with the property that each r-element subset of $\mathcal{S}$ is contained in exactly one block. The matrix representation of this system is a $\binom{n}{r}/\binom{h}{r} \times n$ binary matrix such that each row contains precisely h ones and each column has $\binom{n-1}{r}/\binom{h-1}{r}$ ones. Let $H_{m_0 \times m_0}$ to be a Hadamard matrix of order m_0 [11]. The proposed ETF design in [15] is obtained by substituting non-zero elements of a $S(r, h, n)$ with an arbitrary row of H which contains -1. Therefore, the fingerprint matrix $\mathbf{F}_{ETF}(r, h, n, m_0)$ is an $N \times M$ matrix with entries in $\{-1, 0, +1\}$ where $M = m_0 n$ and $N = \binom{n}{r}/\binom{h}{r}$. It is usually assumed that $n > 8h$.

Next result derives a lower bound on the number of attackers that can break the fingerprinting system given by $\mathbf{F}_{ETF}(r, h, n, m_0)$. The system is broken if the attackers can estimate the host signal $\mathbf{s}$ accurately with high probability.

Theorem 4.1: Let $\boldsymbol{F}_{ETF}(r, h, n, m_0)$ denote the ETF fingerprinting matrix and let K denote the size of coalition set. If

$$K \geq 2 \log 4N/\delta, \tag{11}$$

where $N = \binom{n}{r}/\binom{h}{r}$, then the host signal can be correctly estimated with probability greater than $1 - \delta$.

Proof: Proof is in Appendix. ∎

V. Attack on Random Structure

In this section, we study the performance of the proposed attack on random fingerprints from a finite alphabet.

Random finite fingerprint matrix $\mathbf{F}$ is an N by M matrix whose entries are chosen i.i.d at random from $\Xi = \{\xi_1, ..., \xi_l\}$ with probability $p(f_{i,j} = \xi_k) = p_{\xi_k}$ for $k \in \{1, ..., l\}$. Moreover, we assume that the fingerprints satisfy the energy constraint (3) and without loss of generality, we assume $\mathbb{E}[f_{i,j}] = 0$.

Recall that the attackers might fail to estimate the ith component of the host signal correctly if $\mathcal{A}_i \cap \mathcal{U} = \emptyset$, where $\mathcal{A}_i$ and $\mathcal{U}$ are given by (6) and (5), respectively. Let $\mathbf{b} \in \Xi^K$ and $\tilde{b}_i := b_1 - b_i$; then $\mathcal{B}(\tilde{\mathbf{b}})$, where $\mathcal{B}(\cdot)$ is given by (9), denotes the set of all vectors in Ξ^K that are consistent with the observation vector $\tilde{\mathbf{b}}$.

Attackers estimate $\mathbf{b}$ by $\mathbf{c}^*(\tilde{\mathbf{b}}) = \arg\max_{\mathbf{c} \in \mathcal{B}(\tilde{\mathbf{b}})} p(\mathbf{c})$ that is the most likely vector in $\mathcal{B}(\tilde{\mathbf{b}})$. Since the entries of $\mathbf{F}$ are i.i.d., the probability of error in estimating any component of $\mathbf{s}$ is the same and is given by

$$p(E) = \sum_{\mathbf{b} \in \Xi^K} p(\mathbf{b}) p\left(\mathbf{b} \neq \mathbf{c}^*(\tilde{\mathbf{b}})\right), \tag{12}$$

where $p(\mathbf{b})$ is defined in (10).

Computing (12) in a general setting might be complicated. However, there are some settings in which we can compute $p(E)$ explicitly or at least bound it. The rest of this section focuses on such scenarios.

Definition 5.1: A random fingerprinting matrix $\mathbf{F}_{N \times M}$ with parameters $w \in \mathbb{N}$ and $\mathbf{p} = (p_0, ..., p_w) > 0$ is called symmetric if $f_{i,j}$ are chosen i.i.d. from $\Xi = \{-\frac{w}{z}, ..., -\frac{1}{z}, 0, \frac{1}{z}, ..., \frac{w}{z}\}$, such that $p(f_{i,j} = \frac{k}{z}) = p(f_{i,j} = -\frac{k}{z}) = p_k$ for every $k \in \{0, 1, ..., w\}$.

Note that z is the normalization factor for the fingerprints in order to satisfy (3). Furthermore, $F_{N \times M}$ is called *uniformly symmetric* if $p_k = \dfrac{1}{2w + 1}$ for all k.

Lemma 5.2: The proposed attack succeeds against a uniformly symmetric fingerprinting system with probability greater than $1 - \delta$ as long as the number of colluders, K satisfies

$$K \geq \frac{\log(N/\delta)}{\log(1 + \frac{1}{2w})}.$$

Proof: Proof is in Appendix. ∎

In general, the source distribution $\mathbf{p}_\Xi = (p_1, ..., p_l)$ are not known to the attackers. However, in this case, attackers can estimate the source distribution, for example, using an empirical estimator over those rows of $\mathbf{A}$ that have non-empty intersection with the set $\mathcal{U}$. Namely, they find $\mathcal{I} := \{i : \mathcal{A}_i \cap \mathcal{U} \neq \emptyset\}$. As discussed in previous section, the attackers can determine $(f_{i,1}, ..., f_{i,K})$ correctly for every $i \in \mathcal{I}$. Because the entries of $\mathbf{F}$ are chosen i.i.d., the attackers can estimate $\mathbf{p}_\Xi$ over the set $\{f_{i,j} : i \in \mathcal{I}, 1 \leq j \leq K\}$.

A. Random Ternary Fingerprints

A random ternary fingerprint (RTF) is an $N \times M$ random symmetric fingerprinting system with parameter $w = 1$, and $\mathbf{p} = (1 - 2p, p)$. Next result gives a lower bound on the number of attackers that can detect the host signal with high probability.

Lemma 5.3: An application of the proposed attack on an RTF system succeeds with probability greater than $1 - \delta$ as long as

$$K \geq \frac{\log(\frac{2N}{\delta})}{\log 4/3}, \tag{13}$$

where K is the number of colluders.

Proof: Proof is in Appendix. ∎

In Section VIII we will see the performance of the proposed attack on RTFs for different values of N and M through simulations.

VI. Attack on Probabilistic Fingerprint Codes

This section studies the performance of the proposed attack on binary random fingerprinting codes for M users with length N. Specifically, we consider two such fingerprinting structures. In the first structure which is proposed in [17], each row of the fingerprinting matrix $\mathbf{F}$ is generated independently from a fixed distribution. In the second structure, each user (column), independent of other users, generates its own fingerprints.

A. Row-Wise Generated (Tardos Code)

In this section we perform our attack on a probabilistic structure known as optimal probabilistic fingerprint (OPF) Codes, introduced in [17]. These codes are an improvement of the codes proposed by Boneh-Shaw in [18]. The OPF codes are ϵ-secure (see the following definition) against K pirates and have length $O(K^2 \log 1/\epsilon)$.

Definition 6.1: [17] Let σ be an algorithm that takes a string $\mathbf{y} \in \Xi^N$ (forged copy by the colluders) as input, and produces a subset $\sigma(\mathbf{y}) \subseteq \{1, ..., M\}$ (the set of accused users). For $\emptyset \neq \mathcal{K} \subseteq \{1, ..., M\}$ a $\mathcal{K}$-strategy is an algorithm ρ that takes the sub-matrix of $\mathbf{F}_{N\times M}$ formed by the rows with indices in $\mathcal{K}$ as input, and generate a forgery $\mathbf{y} \in \Xi^N$ as output and satisfies the *marking condition* that for all positions $1 \leq i \leq N$, if all the values $F_{i,j}$ for $j \in \mathcal{K}$ agree with some letter $\xi \in \Xi$, then $y_i = \xi$. A code is called ϵ-secure against coalition of size K, if for any $\mathcal{K} \subseteq \{1, ..., M\}$ with $|\mathcal{K}| \leq K$ and for any $\mathcal{K}$-strategy ρ, the error probability

$$p\left(\sigma(\rho(\mathbf{F}_{\mathcal{K}})) = \emptyset \ or \ \sigma(\rho(\mathbf{F}_{\mathcal{K}})) \nsubseteq \mathcal{K}\right) \leq \epsilon.$$

The code is constructed as follows [17]: Let $c = \lceil \log(1/\epsilon) \rceil$. The fingerprinting matrix $\mathbf{F}_{OPF}$ is of size $(100K^2c =)N \times M$. Let ϱ_i be independent, identically distributed random variable from $[1/(300K), 1 - 1/(300K)]$ for all $i \in \{1, ..., N\}$, where $\varrho_i = \sin^2(r_i)$ and $r_i \sim \text{Uniform}[t, \pi/2 - t]$ with $0 < t < \pi/4$, $\sin^2(t) = 1/(300K)$.

$\mathbf{F}_{OPF}$ is constructed by selecting entry $f_{i,j}$ independently from $\{0, 1\}$ with $p(f_{i,j} = 1) = \varrho_i$ for $i \in \{1, ..., N\}$.

Notice that the OPF code is similar to the previous structure, except that, the fingerprinting matrix is constructed row-wise (row i with probability ϱ_i) instead of column-wise (column j with probability p_j).

The detection procedure for this structure is as follows [17]. The detector defines matrix $\mathbf{U}_{N\times M}$ such that

$$u_{i,j} = \begin{cases} \sqrt{\frac{1-\varrho_i}{\varrho_i}}, & \text{if } f_{i,j} = 1 \\ -\sqrt{\frac{\varrho_i}{1-\varrho_i}}, & \text{if } f_{i,j} = 0 \end{cases}$$

and $Z := 20cK$. Then it accuses user j whenever $(\mathbf{y}^T\mathbf{U})_j > Z$, where $\mathbf{y}$ is the forgery.

Similar to the previous section, consider the coalition $\mathcal{K} = \{1, ..., K\}$ who performs the proposed attack in Section III on this structure. The colluders estimate the probabilities by

$$\widehat{\varrho}_i := \frac{m_i}{|\mathcal{K}|}, \ i \in \mathcal{I}, \tag{14}$$

where $m_i := |\{f_{i,j} = 1 : \ j \in \mathcal{K}\}|$. The attacker's estimation of $\mathbf{f}_1$ will be

$$\widehat{f}_{i,1} = \begin{cases} f_{i,1} & \text{if } i \in \mathcal{I}, \\ 1 & \text{otherwise.} \end{cases}$$

Theorem 6.2: Suppose a coalition of size $K \geq 4$ perform the proposed attack on the above fingerprinting system with matrix $\mathbf{F}_{N\times M}$. Then the number of rows that they fail to estimate correctly $(N - |\mathcal{I}|)$ is bounded by $2CN/\sqrt{K}$ with probability at least

$$1 - \frac{\sqrt{K} - C}{NC},$$

where $1 < C < 1.3$ is a constant.

Proof: Proof is in Appendix. ∎

Colluders use Algorithm 1 to forge $\mathbf{y}$ and hide their tracks as follows: They estimate the signal $\hat{\mathbf{s}}$ and its Fourier transform $\mathcal{F}[\hat{\mathbf{s}}]$. Then add Gaussian and uniform noises to the magnitude and the phase, respectively. They add a random vector $c\mathbf{w}$ in the time domain to form the final forgery. Vector $\mathbf{w}$ is constructed by setting the elements at the indices that the attackers can detect (i.e., $\mathcal{I}$) to zero while choosing a random element from $\{-1, 0, 1\}$ for the other elements.

Note that $\mathcal{F}[\cdot]$ in this algorithm denotes the Fourier transform and $\measuredangle z$ denotes the phase of a complex number z.

Algorithm 1

Input : $\mathcal{I}, \mathbf{q}_1, \widehat{\mathbf{f}}_1, (\sigma_0^2, c_1, c_2) > 0$.
Output : $\mathbf{y}$.
for $i = 1, ..., N$ **do**
  If $i \in \mathcal{I}$; $w_i \longleftarrow 0$.
  If $i \notin \mathcal{I}$; $w_i \longleftarrow \{0, 1, -1\}$ $w.p.$ $\{1/2, 1/4, 1/4\}$.
end for
$\widehat{\mathbf{s}} \longleftarrow \mathbf{q}_1 - \widehat{\mathbf{f}}_1$
$|\mathcal{F}[\widehat{\mathbf{s}}]| \longleftarrow |\mathcal{F}[\widehat{\mathbf{s}}]| + a_1$; $a_1 \sim \mathcal{N}(0, \sigma_0^2)$.
$\measuredangle\mathcal{F}[\widehat{\mathbf{s}}] \longleftarrow \measuredangle\mathcal{F}[\widehat{\mathbf{s}}] + a_2$; $a_2 \sim \text{Uniform}[0, \pi/(2Kc_1)]$.
$\mathbf{y} \longleftarrow \mathcal{F}^{-1}[\mathcal{F}[\widehat{\mathbf{s}}]] + c_2\mathbf{w}$.

Section VIII-E demonstrates the performance of this attack on a OPF code.

B. Column-Wise Generated Fingerprints

Each entry of the fingerprinting matrix $\mathbf{F}_{N\times M}$ is chosen independently from $\{0, 1\}$ with $p(f_{i,j} = 1) = p_j, j = 1, ..., M$, where $p_j = \sin^2(r_j)$ is selected by picking uniformly at random the value $r_j \in [t, \pi/2 - t]$ for some fixed $0 < t < \pi/4$, $\sin^2(t) = a > 0$. Note that in this set up, since the alphabet set is $\{0, 1\}$, we have $\mathcal{U} = \{-1, 1\}$.

Suppose a coalition w.l.o.g. $\mathcal{K} = \{1, ..., K\}$ performs the proposed attack in section III on this structure. Let $\mathcal{I} \subseteq \{1, ..., N\}$ be the set of all indices that the colluders could estimate their fingerprints correctly. For this structure, $\mathcal{I}$ is the index set of those rows in $\mathbf{A} := [\mathbf{q}_1 - \mathbf{q}_1, ..., \mathbf{q}_1 - \mathbf{q}_K]$ that have at least one non-zero entry.

Next step, the colluders estimate their corresponding p_js using empirical estimation over rows with index set $\mathcal{I}$, namely

$$\widehat{p}_j := \frac{n_j}{|\mathcal{I}|}, \quad j \in \mathcal{K}, \tag{15}$$

where $n_j = |\{f_{i,j} = 1 : i \in \mathcal{I}\}|$. Theorem 6.3 describes how accurate this estimation will be. Finally, the attackers use the above empirical distributions to decide about those rows with indices belonging to $\{1, ..., N\} \setminus \mathcal{I}$. To do so, they compute

$$p_{tot} := \frac{1}{K} \sum_{j \in \mathcal{K}} \widehat{p}_j.$$

Then for $i \in \{1, ..., N\} \setminus \mathcal{I}$ and $j \in \mathcal{K}$,

$$\widehat{f}_{i,j} = \begin{cases} 1, & \text{if } p_{tot} > 0.5 + \tau \\ 0, & \text{if } p_{tot} < 0.5 - \tau \end{cases}$$

In case of $0.5-\tau \leq p_{tot} \leq 0.5+\tau$, they pick n_* indices out of $\{1,...,N\} \setminus \mathcal{I}$ randomly and set them to be one and the rest to be zero, where n_* is chosen such that the empirical distributions do not vary, namely

$$\left|\min_j \widehat{p}_j - \frac{n_* + |\mathcal{I}| \min_j \widehat{p}_j}{N}\right| = \left|\max_j \widehat{p}_j - \frac{n_* + |\mathcal{I}| \max_j \widehat{p}_j}{N}\right|.$$

Equivalently,

$$n_* := \left\lfloor \frac{\min_j \widehat{p}_j + \max_j \widehat{p}_j}{2}(N - |\mathcal{I}|) \right\rfloor.$$

The final forgery in this attack will be

$$\mathbf{y} = \mathbf{q}_1 - \widehat{\mathbf{f}}_1.$$

Theorem 6.3: Consider a probabilistic fingerprinting code $\mathbf{F}$ of size $N \times M$ as it is described above and let $\mathcal{K}$ be a coalition of size K. Then

- the number of rows that they fail to estimate correctly $(N - |\mathcal{I}|)$ is bounded by N/K with probability at least

$$1 - 12K^2\left(\frac{3}{8}\right)^K - \frac{8K^2}{N2^K},\ K > 6.$$

- $\mathbb{E}[||\mathbf{f}_1 - \widehat{\mathbf{f}}_1||^2] \leq N/2^{K-1}$.
- with probability at least $1 - \frac{2K}{N^{(2-1/2^{K-2})}}$, we have

$$\max_{j \in \mathcal{K}} |p_j - \widehat{p}_j| < \sqrt{\frac{\log N}{N}}.$$

Proof: Proof is in Appendix. ■

VII. Attack on random Gaussian structure

In this section, we generalize the proposed attack for the case that the alphabet set is continuous.

Suppose the entries of the fingerprinting matrix $\mathbf{F}_{N\times M}$ are chosen i.i.d. at random from a normal distribution $\mathcal{N}(0,\sigma^2)$, such that (3) holds i.e. $\sigma^2 = 1/N$. The main idea here is to convert this fingerprinting structure into a finite alphabet structure by quantizing the support set of the fingerprints and then apply the attack strategy of Section III.

Using the Chebyshev's inequality, the attackers can assume that there exists a $\xi > 0$ such that the fingerprinting entries, $\{f_{k,m}\}$s are chosen only from a bounded interval $[-\xi, \xi]$ with high probability. This is a good assumption since for any $\epsilon > 0$, there exists a $\xi > 0$, such that

$$p\left(|f_{i,j}| \geq \xi\right) \leq \epsilon.$$

Furthermore, the attackers discretize the values in this interval into $2w, w \in \mathbb{N}$ discrete values $\{\widehat{\xi}_i\}_{i=1}^{2w}$ as follows: they assign $\widehat{\xi}_i = -\xi + \frac{\xi}{2w}(2i-1)$ to all the values that belong to interval $[-\xi + \frac{\xi}{w}(i-1), -\xi + \frac{\xi}{w}i)$ for some $i \in \{1,...,2w\}$ and use $\{\widehat{\xi}_i\}_{i=1}^{2w}$ as their alphabet set. Their next step is to represent the entries of the matrix $\mathbf{A}$ introduced in Section III (whose ith column is given by $\mathbf{q}_1 - \mathbf{q}_i$) using the new alphabet set i.e., $\{\widehat{\xi}_i\}_{i=1}^{2w}$. This can be done using the fact that each entry of the matrix $\mathbf{A}$ will fall into $[\frac{\xi}{w}(i-j-1), \frac{\xi}{w}(i-j+1))$, for some $i, j \in \{1,...,2w\}$ with probability at least $(1-\epsilon)^2$. The attackers estimate $\mathbf{A} = [a_{k,m}]$ by $\widehat{\mathbf{A}}_\alpha = [\widehat{a}_{k,m}]$, where

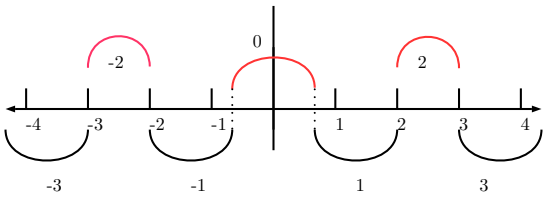

Fig. 1: Quantization rule for the entries of the matrix $\mathbf{A}$, when $\xi = 2$, $w = 2$, and $\alpha = 1.4$.

$$\widehat{a}_{k,m} = \begin{cases} -\frac{\xi}{w}(2w-1), & \text{if } a_{k,m} < -2\xi \\ \frac{\xi}{w}(i-j), & \text{if } a_{k,m} \in [\frac{\xi}{w}(i-j-1), \frac{\xi}{w}(i-j)) \\ & \quad for\ i < j \\ -\frac{\xi}{w}, & \text{if } a_{k,m} \in [-\frac{\xi}{w}, -\alpha\frac{\xi}{2w}) \\ 0, & \text{if } |a_{k,m}| \leq \alpha\frac{\xi}{2w} \\ \frac{\xi}{w}, & \text{if } a_{k,m} \in (\alpha\frac{\xi}{2w}, \frac{\xi}{w}] \\ \frac{\xi}{w}(i-j), & \text{if } a_{k,m} \in (\frac{\xi}{w}(i-j), \frac{\xi}{w}(i-j+1)] \\ & \quad for\ i > j \\ \frac{\xi}{w}(2w-1), & \text{if } a_{k,m} > 2\xi \end{cases} \tag{16}$$

for some $i, j \in \{1,...,2w\}$ and $0 < \alpha \leq 2$. Next we provide an example for the case when $\xi = 2$ and $w = 2$ to motivate our quantization choices.

Example 2: Let $\xi = 2$ and $w = 2$. The attacker's alphabet will be $\{-3/2, -1/2, 1/2, 3/2\}$. If $a_{k,m} := f_{k,1} - f_{k,m} \in [-4,-3)$, then with probability at least $(1-\epsilon)^2$, we have $-2 < f_{k,1} < -1$ and $1 < f_{k,m} < 2$. In other word, $\hat{f}_{k,1} = -3/2$ and $\hat{f}_{k,m} = 1/2$, equivalently, $\widehat{a}_{k,m} = -3/2 - 1/2 = -3$. Figure 1 illustrates this quantization.
Now, suppose $a_{k,m}$ belongs to $(-3,-2)$. In this case, there are three possible events: $E_1 := \{-2 < f_{k,1} < -1, 1 < f_{k,m} < 2\}$, $E_2 := \{-2 < f_{k,1} < -1, 0 < f_{k,m} < 1\}$, or $E_3 := \{-1 < f_{k,1} < 0, 1 < f_{k,m} < 2\}$. If E_1 is the true event, then $\widehat{a}_{k,m} = -3/2 - 1/2 = -3$. However, if either E_2 or E_3 is the true event, then $\widehat{a}_{k,m} = -1/2 - 3/2 = -2$. Since $p(E_2 \cup E_3) > p(E_1)$, we encode $\widehat{a}_{k,m} = -2$. The quantization rule given by (16) is obtained similarly, except for the case that $a_{k,m} \in [-1,1)$. In this case, we consider a small gap interval controlled by α. When $\alpha = 0$, $\widehat{a}_{k,m} = 0$ is neglected and when $\alpha = 2$, the whole interval $[-1,1]$ is quantized into 0.

Without loss of generality, suppose the first K users are involved in the attack and $\widehat{\mathbf{f}}_i$ is the attacker's estimate of $\mathbf{f}_i$ for $1 \leq i \leq K$. Then their final forgery is

$$\widehat{\mathbf{y}} = \frac{1}{K}\sum_{i=1}^{K}(\mathbf{q}_i - \widehat{\mathbf{f}}_i) + \mathbf{w}, \tag{17}$$

where $\mathbf{w}$ is a noise vector with mean zero and variance σ_0^2 introduced by the attackers in order to hide their tracks. Note that one can easily find an upper bound on the expected

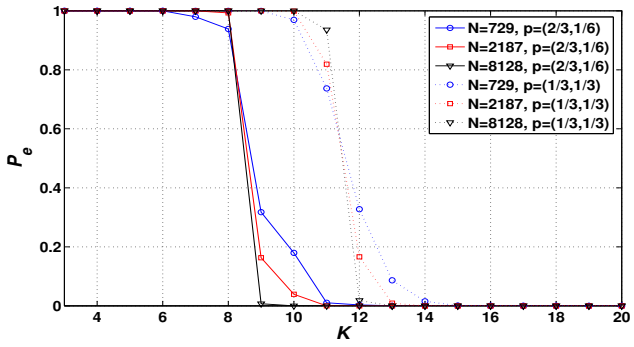

Fig. 2: Probability of estimating at least 1% of the host signal incorrectly as a function of the collusion size K for RTF structure.

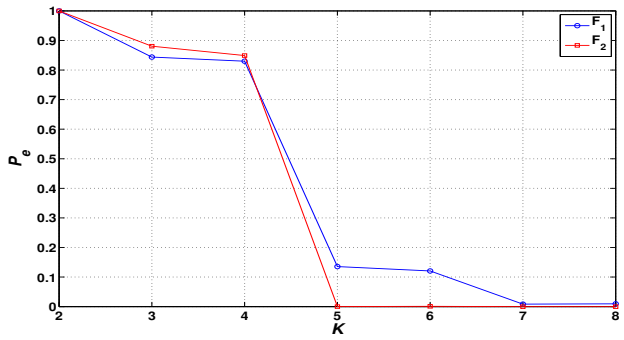

Fig. 3: Probability of $||\mathbf{s}-\hat{\mathbf{s}}|| > 0$ as a function of collusion size K for ETF structure.

estimation error of this attack as follows:

$$\mathbb{E}\left[||\mathbf{f}_1-\widehat{\mathbf{f}}_1||^2\right]^{1/2} \leq \sqrt{N\max_j \mathbb{E}[|f_{1,j}-\widehat{f}_{1,j}|^2]} \leq$$
$$\sqrt{N}\left(\int_0^\infty \frac{(x+(2w-1)\xi/2w)^2}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}dx\right)^{1/2} \leq$$
$$\sqrt{\frac{N}{2}}\left(\sigma+\frac{(2w-1)\xi}{2w}\right).$$

Since $\sigma = 1/\sqrt{N}$, by letting $\xi = O(1/\sqrt{N})$, the above error is bounded as N grows. Later in Section VIII, we study the performance of this attack when $w = 2$ and compare its performance with the uniform linear averaging attack which was proved to outperform order-statistic collusion attacks on random Gaussian fingerprints when the colluders are subject to a mean-squared distortion constraint [6]. The simulation result demonstrates the excellent performance of the proposed algorithm even in continuous case.

VIII. Simulation Results

Herein, we simulate the performance of our attack on both random and deterministic structures.

A. Random Ternary Fingerprints

We attacked two RTF systems with probability vectors $\mathbf{p}_1 = (2/3, 1/6)$ and $\mathbf{p}_2 = (1/3, 1/3)$ each for three different dimensions $(N, M) \in \{(729, 2016), (2187, 8128), (8128, 16384)\}$. Figure 2 depicts the probability of failing to estimate the signal $\mathbf{s}$ correctly as a function of the number of colluders for the random ternary structures. The failure event is defined as estimating at least 1% of the host signal incorrectly. As Figures 2 illustrates that the number of colluders necessary to estimate the host signal in RTF system is much less than what Lemma 5.3 suggests. This confirms the effectiveness of our proposed attack in practice beyond the theoretical limit.

B. ETF Fingerprints

We attacked two ETF systems with matrices $\mathbf{F}_1 = \mathbf{F}_{ETF}(2, 7, 91, 16)$ and $\mathbf{F}_2 = \mathbf{F}_{ETF}(2, 2, 2^7, 128)$, respectively. Figure 3 depicts the probability of failure event for these two ETF structures, when the failure event is defined as at least estimating one coordinate of the host signal $\mathbf{s}$ incorrectly. Again the number of users that succeed at estimating the host signal is much less than what the bounds in Theorem 4.1 suggests.

C. Gaussian Fingerprints

We also attacked two random Gaussian fingerprinting systems with dimensions $(N, M) \in \{(100, 250), (1000, 2070)\}$. The attackers parameters were set at $(\xi, w, \alpha) = (0.12, 2, 1)$ and the forgery was done according to (17). Figure 4 depicts $||\mathbf{f}_{i_1} - \widehat{\mathbf{f}}_{i_1}||$, the estimation error of the attackers i, as a function of the collusion size K.

We compared the performance of our attack against the uniform linear averaging attack, the optimal attack in the class of order-statistic collusion attacks on random Gaussian fingerprints when the colluders are subject to a mean-squared distortion constraint [6]. The colluders in that attack create their forgery by uniformly averaging their copies and adding an i.i.d. Gaussian noise. We evaluated performance of both attacks for a focused fingerprint detector, which performs a binary hypothesis test for each user to decide whether that particular user is guilty [10].

Figures 5 and 6 show the probability of detecting at least one colluder as a function of the number of colluders K using the focused detector.

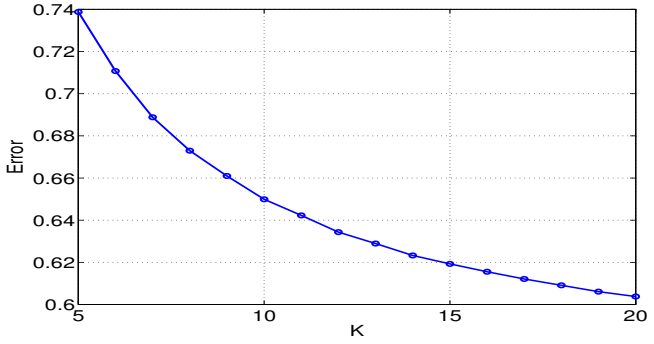

Fig. 4: Estimation error $||\mathbf{f}_{i_1} - \widehat{\mathbf{f}}_{i_1}||$ for random Gaussian fingerprints ($(N, M) = (1000, 2070)$).

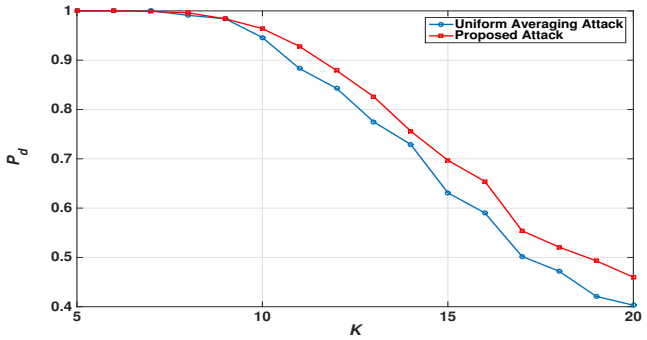

Fig. 5: A comparison of probability of detecting at least one colluder for our proposed attack vs. uniform linear averaging attack as a function of number of colluders for random Gaussian fingerprints ($(N, M) = (1000, 2070)$).

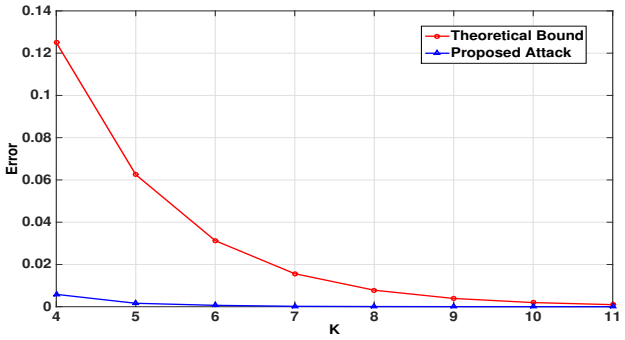

Fig. 7: Worst estimation error rate of the fingerprints and its corresponding theoretical bound for the CWC fingerprinting structure.

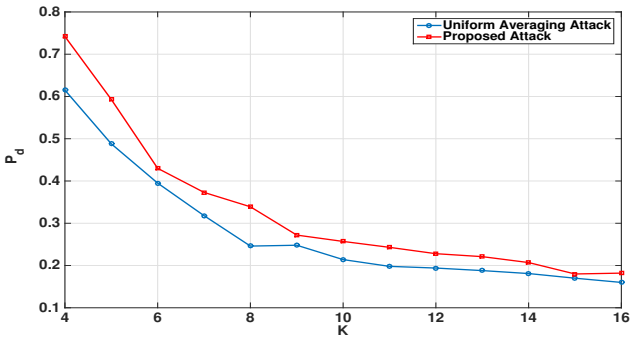

Fig. 6: A comparison of probability of detecting at least one colluder for our proposed attack vs. uniform linear averaging attack as a function of number of colluders for random Gaussian fingerprints ($(N, M) = (100, 250)$).

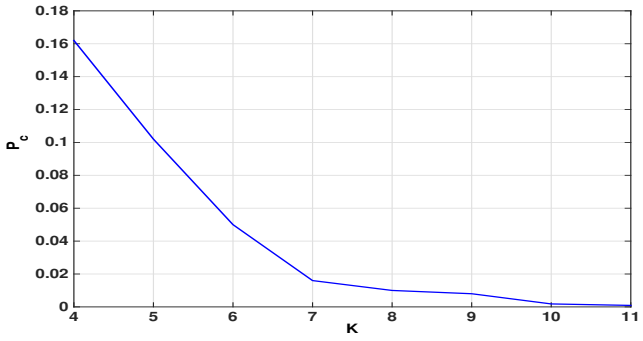

Fig. 8: A plot of the probability of catching at least one colluder P_c as a function of the number of colluders for the CWC structure.

D. Tardos-like (Column-Wise Random) fingerprints

The performance of the proposed attack against the fingerprinting code proposed in Section VI-B is studied in this section. We generated a fingerprinting code with parameters $(N, M, t') = (5600, 2100, \pi/1000)$. $K \in \{4, ..., 11\}$ number of colluders were selected randomly 500 times to arrange the attack against this code.

$K \in \{4, ..., 11\}$ number of colluders were selected randomly 500 times to carryout the attack. Figure 7 demonstrates the worst estimation error rate incurred by an attacker, i.e., $\max_{j\in\mathcal{K}} ||\mathbf{f}_j - \widehat{\mathbf{f}}_j||^2/N$ along with the theoretical bound introduced in Theorem 6.3 as a function of number of colluders. The plot suggests that in practice the performance of our attack is considerably better than predicted by the theoretical bound.

Additionally, we evaluated the performance of our attack against the focused detector which aims to identify at least one of the colluders. We selected K attackers and performed the attack 500 times at random. Figure 8 illustrates the probability of catching at least one colluder as a function of number of colluders. In the figure, probability of catching P_c is calculated as

$$P_c := \frac{\#\text{trials that at least one colluders is detected}}{\#\text{total trials}}. \quad (18)$$

E. Row-Wise Generated Fingerprints (Tardos Code)

We simulated our proposed attack against the optimal probabilistic fingerprints (OPF), known as Tardos codes, introduced in Section VI-A. We considered a length $N = 7500$ OPF code designed with parameters $(k, \epsilon) = (5, 0.1)$ for $M = 1500$ number of users. Our first measure of success of the attack is the size of the set of fingerprint indices the colluders detected correctly, $|\mathcal{I}|$. Figure 9 demonstrates the rate $|\mathcal{I}|/N$ as a function of the coalition size for $K = \{2, 3, 4, 5\}$.

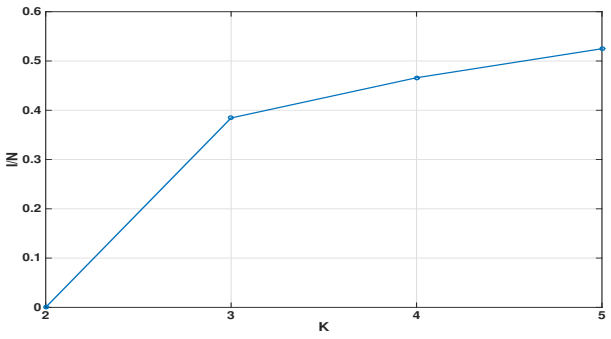

Fig. 9: A plot of $|\mathcal{I}|/N$ as a function of the number of colluders for the Tardos code with parameter $(k, \epsilon) = (5, 0.1)$.

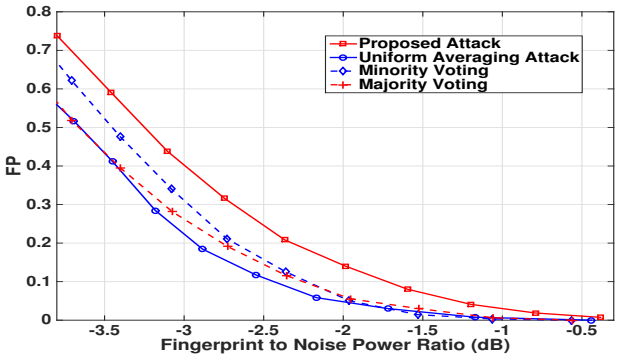

Fig. 10: Comparison of the fales positive (FP) probabilities of different attacks against the Tardos code with parameter $(k, \epsilon) = (5, 0.1)$.

Next, we selected $K = 5$ users randomly and attacked the fingerprinting scheme using the forgery described by Algorithm 1 for 3000 trials. We compared the performance of our attack against the uniform averaging, minority voting, and majority voting attacks [12], [14]. In minority voting attack, the attackers output a symbol that occurs less often than all the others in their copies. Similarly, in majority voting, the attackers output the most frequent symbol [13]. All attacks were evaluated against the optimal accusation detector of Tardos [17], which was described in Section VI-A. We computed the false positive (FP) probability, i.e., the probability of accusing at least one innocent user, as a function of the ratio of the fingerprint power to the noise power introduced by the colluders (Ratio) for all 3 attacks. More precisely,

$$\text{Fingerprint to Noise Power Ratio (FNPR)} := 20\log_{10}\frac{||\mathbf{f}_1||}{||\mathbf{y}-\mathbf{s}||},$$

where $\mathbf{y}$ is the forgery. As Figure 10 illustrates, the proposed attack outperforms the others since for the same number of colluders and the same FNPR, it results in a higher false positive.

IX. Conclusion

We develop a new attack on fingerprinting system with finite alphabet fingerprints. We derive theoretical bounds on the number of colluders that can defeat the fingerprinting system. Through simulation, we show that a much smaller number of attackers can indeed defeat the system in practice. Our analysis suggests that the designer of a random fingerprinting system with parameters $(\Xi, \mathbf{p}_\Xi)$ should select the fingerprints from an alphabet Ξ whose associated uniquely decodable set, $\mathcal{U}$, has a small cardinality. Moreover, the fingerprinter should pick a source distribution $\mathbf{p}_\Xi$ that maximizes the error probability (12) subject to the energy constraint (3).

Acknowledgment

X. Appendix

Proof of Theorem 4.1

Assume K colluders $\{i_1, ..., i_K\}$ take part in the forgery. Let E_i to be the event that the attackers estimate s_i incorrectly, i.e., $E_i := \{\hat{s}_i \neq s_i\}$. Since the attackers know that the fingerprint matrix is sparse, when they cannot estimate f_{i,i_1}[1] uniquely (i.e., $\mathcal{A}_i \cap \mathcal{U} = \emptyset$), the performance of the proposed attack is lower bounded by the following attack. They estimate f_{i,i_1} with an element of $\mathcal{B}\left((f_{i,i_1} - f_{i,i_1}, ..., f_{i,i_1} - f_{i,i_K})\right)$ with the maximum number of zeros. The error event E_i only occurs when vector $(f_{i,i_1}, ..., f_{i,i_K})$ both does not contain the pair $(-1, +1)$ and the number of its non-zero elements exceeds the number of its zeros. Since each row of $\mathbf{F}_{ETF}(r, h, n, m_0)$ has exactly hm_0 non-zero elements, and each row of the corresponding Hadamard matrix has equal number of $+1$ and -1, each row of $\mathbf{F}_{ETF}(r, h, n, m_0)$ has exactly $hm_0/2$, -1s. Therefore, we can upper bound the error event as

$$p(E_i) \leq 2\sum_{i\leq\lfloor K/2\rfloor}\binom{M-hm_0}{i}\binom{hm_0/2}{K-i}/\binom{M}{K}. \quad (19)$$

Next we introduce an upper bound on the above error probability, first we show that for $K \geq 2$ and $n > 8h$,

$$\sum_{i\leq\lfloor K/2\rfloor}\binom{M-hm_0}{i}\binom{hm_0/2}{K-i} \quad (20)$$
$$\leq 2\binom{M-hm_0}{\lfloor K/2\rfloor}\binom{hm_0/2}{\lfloor (K+1)/2\rfloor}.$$

The proof is by induction on K. The case $K = 2$ is trivial since $M > 8hm_0$. Suppose that the inequality holds for $K < k$. In order to prove the inequality for $K = k$ we consider two cases: k is an even ($k = 2q$) or an odd number ($k = 2q+1$). Here we only prove the even case since the proof for the odd case is similar. If $k = 2q$, we need to show

$$\sum_{i\leq q-1}\binom{M-hm_0}{i}\binom{hm_0/2}{2q-i} \leq \binom{M-hm_0}{q}\binom{hm_0/2}{q}.$$

This is true because

$$\sum_{i\leq q-1}\binom{M-hm_0}{i}\binom{hm_0/2}{2q-i}$$
$$= \sum_{i\leq q-1}\binom{M-hm_0}{i}\binom{hm_0/2}{2q-2-i}\frac{hm_0/2-2q+i+1}{2q-i-1}$$
$$\times\frac{hm_0/2-2q+i+2}{2q-i}$$

$$\leq \sum_{i\leq q-1}\binom{M-hm_0}{i}\binom{hm_0/2}{2q-2-i}\frac{hm_0/2-q}{q}$$
$$\times\frac{hm_0/2-q+1}{q+1}$$
$$\leq 2\binom{M-hm_0}{q-1}\binom{hm_0/2}{q-1}\frac{(hm_0/2-q)(hm_0/2-q+1)}{(q+1)q}$$
$$= 2\binom{M-hm_0}{q}\binom{hm_0/2}{q}\frac{(hm_0/2-q)q}{(q+1)(M-hm_0-q+1)}$$
$$\leq \binom{M-hm_0}{q}\binom{hm_0/2}{q}$$

The first inequality is due to the fact that $g(x) := \frac{a-b+x}{b-x+1}$ is an increasing function of x if $a+1 > 0$. The second inequality

[1]Note that for the attackers, estimating at least one of the fingerprints is equivalent to estimating all of their fingerprints.

is because of the induction hypothesis, and the last inequality is valid since $q/(q+1) \leq 1$ and because of the assumption that $M \geq 8hm_0$. Substituting (20) into (19) implies

$$p(E_i) \leq 4\frac{\binom{M-hm_0}{\lfloor K/2 \rfloor}\binom{hm_0/2}{\lfloor (K+1)/2 \rfloor}}{\binom{M}{K}}$$
$$= 4\binom{K}{\lfloor (K+1)/2 \rfloor}\prod_{j=1}^{\lfloor (K+1)/2 \rfloor}\left(1-\frac{M-\lfloor K/2 \rfloor - hm_0/2}{M-K+j}\right)$$
$$\times \prod_{j=1}^{\lfloor K/2 \rfloor}\left(1-\frac{hm_0}{M-\lfloor K/2 \rfloor + j}\right)$$
$$\leq 4\left(\frac{Ke}{\lfloor K/2 \rfloor}\right)^{\lfloor K/2 \rfloor}\left(\frac{hm_0/2}{M-\lfloor K/2 \rfloor}\right)^{\lfloor (K+1)/2 \rfloor}\left(\frac{M-hm_0}{M}\right)^{\lfloor K/2 \rfloor}.$$

The last inequality is because $\binom{n}{n-k} = \binom{n}{k} \leq (ne/k)^k$, where e is the Euler's constant and the fact that $1-a/(b+x)$ is an increasing function of x when $a > 0, b > 0$. By considering K to be even or odd separately, and the assumption that $M > 8hm_0$ one can show that the last equation in the above inequality can be bounded as

$$4\left(\frac{Ke}{\lfloor K/2 \rfloor}\right)^{\lfloor K/2 \rfloor}\left(\frac{hm_0/2}{M-\lfloor K/2 \rfloor}\right)^{\lfloor (K+1)/2 \rfloor}\left(\frac{M-hm_0}{M}\right)^{\lfloor K/2 \rfloor}$$
$$\leq 4\left(\frac{ehm_0(M-hm_0)}{M(M-K/2)}\right)^{K/2}.$$

Recall that $M = nm_0$. The result follows from applying the union bound argument.

Proof of Lemma 5.2

It is clear that this structure satisfies the zero mean condition. Since the elements of fingerprints are chosen uniformly, all the elements in $\mathcal{B}(\tilde{\mathbf{b}})$ have the same likelihood for every $\mathbf{b} \in \Xi^K$. Hence, the error probability (12) can be written as

$$p(E) = \frac{1}{(2w+1)^K}\sum_{\mathbf{b}\in\Xi^K}(1-\frac{1}{|\mathcal{B}(\tilde{\mathbf{b}})|}).$$

It can be proven that due to the symmetric structure of $\mathbf{F}$, $|\mathcal{B}(\tilde{\mathbf{b}})| = t$ when the elements of $\mathbf{b}$ have maximum distance $(2w+1-t)/z$, i.e., $\max_{i,j}|b_i - b_j| = (2w+1-t)/z$. Moreover, $|\mathcal{B}(\tilde{\mathbf{b}})| \leq 2w+1$. Using these facts we obtain

$$\sum_{\mathbf{b}\in\Xi^K}1-\frac{1}{|\mathcal{B}(\tilde{\mathbf{b}})|} = \sum_{t=1}^{2w+1}(1-\frac{1}{t})|\{\mathbf{b}:|\mathcal{B}(\tilde{\mathbf{b}})|=t\}|$$
$$= 2w + \sum_{t=2}^{2w}\Big((t-1)\Big((2w+2-t)^K - 2(2w+1-t)^K + (2w-t)^K\Big)\Big) = (2w)^K.$$

This implies $p(E) = \left(1-\frac{1}{2w+1}\right)^K$ and by applying the union bound, the result is immediate.

Proof of Lemma 5.3

In this scenario, it is clear that the performance of the proposed attack is lower bounded by an attack in which the attackers choose an element in $\mathcal{B}((f_{i,1}-f_{i,1},...,f_{i,1}-f_{i,K}))$ with maximum number of zeros whenever $\mathcal{A}_i \cap \mathcal{U} = \emptyset$. To analyze this attack, we consider two cases: $0 < p \leq 1/4$ and $1/4 < p < 1/2$. If $0 < p \leq 1/4$, it can be shown that the attackers might fail to learn s_i only when both $\mathcal{A}_i \cap \mathcal{U} = \emptyset$ and the number of non-zero elements exceeds the number of zeros. Let $p(E)$ to be the probability of such event, then

$$p(E) \leq 2\sum_{j\geq\lfloor (K+1)/2 \rfloor}\binom{K}{j}p^j(1-2p)^{K-j}. \quad (21)$$

From application of binomial expansion,

$$(3/2)^K = \sum_{j=0}^{\lfloor (K-1)/2 \rfloor}\binom{K}{j}\frac{1}{2^j} + \sum_{j=0}^{\lfloor K/2 \rfloor}\binom{K}{j+\lfloor (K+1)/2 \rfloor}\frac{1}{2^{j+\lfloor (K+1)/2 \rfloor}},$$

we obtain

$$\sum_{j=0}^{\lfloor K/2 \rfloor}\binom{K}{j+\lfloor (K+1)/2 \rfloor}\frac{1}{2^j} \leq 2^{\lfloor (K+1)/2 \rfloor}(3/2)^K. \quad (22)$$

Rewrite the right hand side of (21) as follows

$$2p^{\lfloor (K+1)/2 \rfloor}(1-2p)^{\lfloor K/2 \rfloor}\sum_{j=0}^{\lfloor K/2 \rfloor}\binom{K}{j+\lfloor (K+1)/2 \rfloor}\left(\frac{p}{1-2p}\right)^j$$
$$\overset{(a)}{\leq} 2p^{\lfloor (K+1)/2 \rfloor}(1-2p)^{\lfloor K/2 \rfloor}\sum_{j=0}^{\lfloor K/2 \rfloor}\binom{K}{j+\lfloor (K+1)/2 \rfloor}\frac{1}{2^j}$$
$$\overset{(b)}{\leq} 2(2p)^{\lfloor (K+1)/2 \rfloor}(1-2p)^{\lfloor K/2 \rfloor}(3/2)^K \leq 2\left(\frac{3}{4}\right)^K.$$

Inequality (a) follows because $p \leq 1/4$, thus $2p \leq 1-2p$. Inequality (b) follows from (22) and the last inequality is due to the fact that $\max\{x^{q+1}(1-x)^q, x^q(1-x)^q\} \leq \frac{1}{4^q}$, when $0 \leq x \leq 1-x,\ q \in \mathbb{N}$.

If $1/4 < p < 1/2$, we assume the attackers make an error any time $\mathcal{A}_i \cap \mathcal{U} = \emptyset$ or equivalently when $(f_{i,1},...,f_{i,K})$ does not include the pair $(-1,1)$. This happens with probability less than $2(1-p)^K - (1-2p)^K$ which is also less than $2\,(3/4)^K$ for $1/4 < p < 1/2$. Hence, $p(E) \leq 2\,(3/4)^K$. The result is immediate from the union bound.

Proof of Theorem 6.2

Let $S = \sum_{i=1}^N \mathbb{I}_{\{\mathbf{A}_i=[0,...,0]\}}$. Hence, S is the number of rows that colluders can not estimate their corresponding fingerprints correctly, and let $p(e_i)$ to be the probability of event $\{\mathbb{I}_{\{\mathbf{A}_i=[0,...,0]\}} = 1\}$. We obtain

$$\mathbb{E}[S] = \mathbb{E}[\mathbb{E}[S|e_i, 1 \leq i \leq N]] \quad (23)$$
$$= \sum_{i=1}^N \mathbb{E}[p(e_i)] = \sum_{i=1}^N \mathbb{E}[\varrho_i^K + (1-\varrho_i)^K]$$
$$= N\left(\int_t^{\pi/2-t}\sin^{2K}(x) + \cos^{2K}(x)dx\right)/(\pi/2-2t).$$

Using the fact that

$$\int \sin^n(x)dx = -\frac{\sin^{n-1}(x)\cos(x)}{n} + \frac{n-1}{n}\int \sin^{n-2}(x)dx,$$

we obtain

$$\int_t^{\pi/2-t}\sin^{2K}(x)dx \quad (24)$$
$$\leq \sin(t) + \left(\tfrac{2K-1}{2K} \times \cdots \times \tfrac{1}{2}\right)(\pi/2-2t).$$

Similarly,

$$\begin{aligned}&\int_t^{\pi/2-t}\cos^{2K}(x)dx \\ &\leq \sin(t)+\left(\tfrac{2K-1}{2K}\times\cdots\times\tfrac{1}{2}\right)(\pi/2-2t).\end{aligned} \tag{25}$$

Substituting (24) and (25) into (23), gives us

$$\begin{aligned}\tfrac{\mathbb{E}[S]}{N} &\leq 2\tfrac{\sin(t)}{\pi/2-2t}+2\left(\tfrac{2K-1}{2K}\times\cdots\times\tfrac{1}{2}\right)\\ &=\tfrac{2}{\sqrt{300K}(\pi/2-2t)}+2\tfrac{\binom{2K}{K}}{4^K}\\ &\overset{a}{\simeq}\tfrac{2}{\sqrt{300K}(\pi/2-2t')}+\tfrac{2}{\sqrt{\pi K}}:=\tfrac{C}{\sqrt{K}}.\end{aligned}$$

Where (a) uses a well known approximation which is $\binom{2K}{K}\simeq\frac{4^K}{\sqrt{\pi K}}$. Furthermore, we obtain

$$\begin{aligned}Var[S]&=\mathbb{E}[\mathbb{E}[S^2|e_i,1\leq i\leq N]]-\mathbb{E}[S]^2\\ &=\textstyle\sum_{i=1}^N\mathbb{E}[p(e_i)]-\mathbb{E}[p(e_i)]^2\\ &\overset{b}{\leq} N(C/\sqrt{K}-C^2/K).\end{aligned}$$

(b) is true since $x-x^2$ is an increasing function for $0<x<0.5$. Finally, Chernoff's bound will establish the result.

Proof of Theorem 6.3

First we compute the mean and variance of S, i.e., the number of rows in $\mathbf{A}$ that attackers can not estimate. S can be written as

$$S=\sum_{i=1}^N\mathbb{I}_{\{\mathbf{A}_i=[0,\ldots,0]\}}$$

where $\mathbb{I}$ is the indicator function. Let e_i to be the event that ith row of $\mathbf{A}$ is entirely zeros. This event happens when either $[f_{i,1},...,f_{i,K}]$ is a vector of all zeros or a vector of all ones. Therefore,

$$p(e_i)=\prod_{j=1}^K p_j+\prod_{j=1}^K(1-p_j), 1\leq i\leq N, \tag{26}$$

where $p_j=\sin^2(r_j)$ and $r_j\sim Uniform(t',\pi/2-t')$. Now one can obtain

$$\begin{aligned}&\mathbb{E}[||\mathbf{f}_1-\widehat{\mathbf{f}}_1||^2]\leq\mathbb{E}[S]\\ &=\mathbb{E}[\mathbb{E}[S|e_i,1\leq i\leq N]]=N\mathbb{E}[p(e_i)]\\ &\overset{a}{=}N\left(\mathbb{E}[p_1]^K+(1-\mathbb{E}[p_1])^K\right)\\ &\overset{b}{=}N/2^{K-1},\end{aligned} \tag{27}$$

where (a) is true since p_js are i.i.d and (b) is because of $\int_t^{\pi/2-t}\sin^2(r)\frac{dr}{\pi/2-2t}=1/2$.

$$\begin{aligned}&\mathbb{E}[S^2]=\mathbb{E}[\mathbb{E}[S^2|e_i,1\leq i\leq N]]\\ &\overset{c}{=}\mathbb{E}[N^2p(e_i)^2+Np(e_i)-Np(e_i)^2]\\ &=(N^2-N)\mathbb{E}[p(e_i)^2]+N\mathbb{E}[p(e_i)]\\ &=(N^2-N)\left(\mathbb{E}[p_1^2]^K+\mathbb{E}[(1-p_1)^2]^K+2\mathbb{E}[p_1(1-p_1)]^K\right)\\ &+N\mathbb{E}[p(e_i)]\overset{d}{\leq}(N^2-N)\left(3\left(\tfrac{3}{8}\right)^K\right)+N/2^{K-1},\end{aligned} \tag{28}$$

where (c) is due to the fact that $S|(e_i,1\leq i\leq N)$ is a binomial distribution with parameter $p(e_i)$ and (d) is because of $5/16<\int_t^{\pi/2-t}\sin^4(r)\frac{dr}{\pi/2-2t}<3/8$. Using (27), (28), and Chernoff's bound, we obtain

$$\begin{aligned}&P\left(S\geq N/K\right)<\\ &P\left(S\geq N/2K+N/2^{K-1}\right)=P\left(|S-N/2^{K-1}|>N/2K\right)\leq\\ &4K^2\left(3\left(\tfrac{3}{8}\right)^K-4\left(\tfrac{2}{8}\right)^K\right)+\tfrac{4K^2}{N}\left(2\left(\tfrac{4}{8}\right)^K-3\left(\tfrac{3}{8}\right)^K\right)\\ &\overset{e}{\leq}12K^2(\frac{3}{8})^K+\frac{8K^2}{N}(\frac{4}{8})^K,\end{aligned}$$

(e) is true because of the fact that $3\cdot3^K-4\cdot2^K<3\cdot3^K$ and $2\cdot4^K-3\cdot3^K<2\cdot4^K$. Therefore, for $K>6$ with probability at least $1-12K^2(3/8)^K-8K^2(1/2)^K/N$, we have $S<N/K$. In this case, using the Hoeffding's inequality and union bound we get

$$P\left(\max_{i\in\mathcal{K}}|p_i-\widehat{p}_i|\geq\rho\right)\leq 2Ke^{-2(N-S)\rho^2}, \tag{29}$$

By choosing $\rho=\sqrt{\frac{\log N}{N}}$, we obtain the result.

REFERENCES

[1] Z. Wang, M. Wu, H. Zhao, W. Trappe, and K. Liu, "Anti-collusion forensics of multimedia fingerprinting using orthogonal modulation," Image Processing, IEEE Transactions on, vol. 14, no. 6, pp. 804-821, 2005.

[2] N. Kiyavash and P. Moulin, "On Optimal Collusion Strategies for Fingerprinting," IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2006.

[3] P. Moulin and N. Kiyavash, "Expurgated Gaussian fingerprinting codes," IEEE International Symposium on Information Theory, 2007.

[4] N. Kiyavash and P. Moulin, "A Framework for Optimizing Nonlinear Collusion Attacks on Fingerprinting Systems," Conference on Information Systems and Science, Princeton, NJ, March 2006.

[5] N. Kiyavash and P. Moulin, and T. Kalker "Regular simplex fingerprints and their optimality properties," IEEE Transactions on Information Forensics and Security, pp. 318-329, Vol. 4, No. 3, 2009.

[6] P. Moulin and N. Kiyavash, "Performance of Random Fingerprinting Codes under Arbitrary Non-linear Attacks," IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2007.

[7] J. Cox, J. Killian, F. T. Leighton, and T. Shamoon, "Secure Spread Spectrum Watermarking for Multimedia," Image Processing, IEEE Transactions on 6.12, pp. 1673-1687, 1997.

[8] "Steiner systems," in La Jolla Covering Repository, 2011, http://www.ccrwest.org/cover/steiner.html

[9] J. Etesami and N. Kiyavash. "A Novel Collusion Attack on Finite Alphabet Digital Fingerprinting Systems." IEEE International Symposium on Information Theory (ISIT), 2014.

[10] D. G. Mixon, C. J. Quinn, N. Kiyavash, and M. Fickus, "Fingerprinting with Equiangular Tight Frames," IEEE Transactions on Information Theory, pp. 1855-1865, Vol. 59, No. 3, 2013.

[11] R. K. Yarlagadda and J. E. Hershey, "Hadamard Matrix Analysis and Synthesis," Boston, 1997, Kluwer. ISBN 0-7923-9826-2.

[12] A. Simone, and B. Skorić,"False Positive probabilities in q-ary Tardos codes: comparison of attacks." Designs, Codes and Cryptography 75.3 (2012): 519-542.

[13] A. Simone, and B Skorić. "Accusation probabilities in Tardos codes: the Gaussian approximation is better than we thought." IACR Cryptology ePrint Archive 2010 (2010): 472.

[14] B. Skorić, S. Katzenbeisser, and M. U. Celik, "Symmetric Tardos fingerprinting codes for arbitrary alphabet sizes." Designs, Codes and Cryptography 46, no. 2 (2008): 137-166.

[15] M. Fickus, D. G. Mixon, and J. Tremain, "Steiner equiangular tight frames," Linear Algebra Appl., vol. 436, no. 5, pp. 1014-1027, 2012.

[16] Z. Shahid, M. Chaumont, and W. Puech. "Spread spectrum-based watermarking for Tardos code-based fingerprinting for H. 264/AVC video." Image Processing (ICIP), 2010 17th IEEE International Conference on. IEEE, 2010.

[17] G. Tardos, "Optimal probabilistic fingerprint codes," Journal of the ACM (JACM) 55, no. 2, 2008.

[18] D. Boneh, J. Shaw, “Collusion-secure fingerprinting for digital data,” IEEE Transactions on Information Theory, v.44 n.5, p.1897-1905, September 1998.